High Performance
Computing &
Big Data Services

hpc.uni.lu
hpc@uni.lu
@ULHPC

HPC & Big Data Analytics in Luxembourg

... at the EuroHPC Horizon

Dr. Sébastien Varrette

University of Luxembourg (UL), Luxembourg

10th Year ISACA Luxembourg Celebration

Sept. 20th, 2018, Luxembourg

Trust in, and value from, information systems
Luxembourg Chapter

Short CV

https://varrette.gforge.uni.lu

- Experienced **Research Scientist** at University of Luxembourg (UL)
  - ↪ 15 y of experience in management & deployment of HPC systems
  - ↪ Ph.D Computer Science w. great distinction (2007, INPG/UL)
  - ↪ **Research interests**:
    - ✓ Security, perf. parallel & distributed computing platforms (HPC/Cloud/IoT)
    - ✓ 590 citations, **4 books**, 80 articles in peer-reviewed journals/conf.

Short CV

https://varrette.gforge.uni.lu

- Experienced **Research Scientist** at University of Luxembourg (UL)
  - ↪ 15 y of experience in management & deployment of HPC systems
  - ↪ Ph.D Computer Science w. great distinction (2007, INPG/UL)
  - ↪ **Research interests**:
    - ✓ Security, perf. parallel & distributed computing platforms (HPC/Cloud/IoT)
    - ✓ 590 citations, **4 books**, 80 articles in peer-reviewed journals/conf.

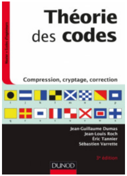

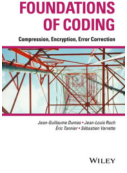

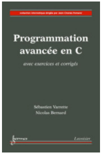

Short CV

https://varrette.gforge.uni.lu

- Experienced **Research Scientist** at University of Luxembourg (UL)
  - ↪ 15 y of experience in management & deployment of HPC systems
  - ↪ Ph.D Computer Science w. great distinction (2007, INPG/UL)
  - ↪ **Research interests**:
    - ✓ Security, perf. parallel & distributed computing platforms (HPC/Cloud/IoT)
    - ✓ 590 citations, **4 books**, 80 articles in peer-reviewed journals/conf.

- **UL HPC Manager** (head: Prof. Bouvry) since 2007
- Management committee member representing Luxembourg:
  - ↪ ETP4HPC, EU COST NESUS, **PRACE** (acting *Advisor*)
- National / EU HPC projects involvement:
  - ↪ **National HPC and Big Data competence center** (MECO)
  - ↪ NVidia Cooperation agreement on AI and HPC (SMC)
  - ↪ EuroHPC / IPCEI on HPC and Big Data (BD) Applications

Summary

Summary

Summary

Why HPC and BD ?

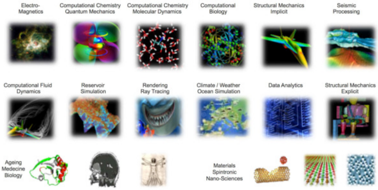

HPC: **H**igh **P**erformance **C**omputing

BD: **B**ig **D**ata

Andy Grant, Head of Big Data and HPC, Atos UK&I

To out-compete you must out-compute

Increasing competition, heightened customer expectations and shortening product development cycles are forcing the pace of acceleration across all industries.

Why HPC and BD ?

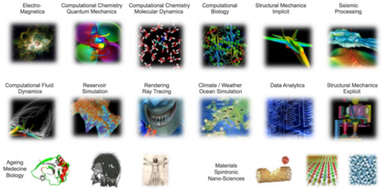

HPC: **H**igh **P**erformance **C**omputing
BD: **B**ig **D**ata

- Essential tools for **Science, Society and Industry**
  - ↪ All scientific disciplines are becoming computational today
    - ✓ requires very high computing power, handles **huge** volumes of data
- **Industry, SMEs** increasingly relying on HPC
  - ↪ to invent innovative solutions
  - ↪ ... while reducing cost & decreasing time to market

Andy Grant, Head of Big Data and HPC, Atos UK&I

To out-compete you must out-compute

Increasing competition, heightened customer expectations and shortening product development cycles are forcing the pace of acceleration across all industries.

Why HPC and BD ?

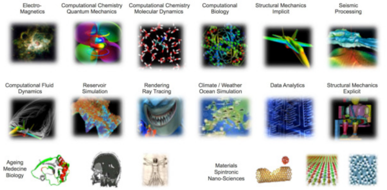

HPC: **H**igh **P**erformance **C**omputing
BD: **B**ig **D**ata

- Essential tools for **Science, Society and Industry**
  - ↪ All scientific disciplines are becoming computational today
    - ✓ requires very high computing power, handles **huge** volumes of data
- **Industry, SMEs** increasingly relying on HPC
  - ↪ to invent innovative solutions
  - ↪ . . . while reducing cost & decreasing time to market
- HPC = **global race** (strategic priority) - EU takes up the challenge:
  - ↪ PRACE / EuroHPC / IPCEI on HPC and Big Data (BD) Applications

Andy Grant, Head of Big Data and HPC, Atos UK&I

To out-compete you must out-compute

Increasing competition, heightened customer expectations and shortening product development cycles are forcing the pace of acceleration across all industries.

Prerequisites: Metrics

- **HPC**: **H**igh **P**erformance **C**omputing — **BD**: **B**ig **D**ata

Main HPC/BD Performance Metrics

- **Computing Capacity**: often measured in **flops** (or **flop/s**)
  - ↪ **Fl**oating **p**oint **o**perations per **s**econds (often in DP)
  - ↪ **GFlops** $= 10^9$ **TFlops** $= 10^{12}$ **PFlops** $= 10^{15}$ **EFlops** $= 10^{18}$

Prerequisites: Metrics

- **HPC**: **H**igh **P**erformance **C**omputing **BD**: **B**ig **D**ata

Main HPC/BD Performance Metrics

- **Computing Capacity**: often measured in **flops** (or **flop/s**)
  - ↪ **Fl**oating **p**oint **o**perations per **s**econds (often in DP)
  - ↪ **GFlops** $= 10^9$ **TFlops** $= 10^{12}$ **PFlops** $= 10^{15}$ **EFlops** $= 10^{18}$
- **Storage Capacity**: measured in multiples of **bytes** = 8 **bits**
  - ↪ **GB** $= 10^9$ bytes **TB** $= 10^{12}$ **PB** $= 10^{15}$ **EB** $= 10^{18}$
  - ↪ **GiB** $= 1024^3$ bytes **TiB** $= 1024^4$ **PiB** $= 1024^5$ **EiB** $= 1024^6$
- Transfer rate on a medium measured in **Mb/s** or **MB/s**
- Other metrics: Sequential vs Random **R/W speed**, **IOPS** . . .

Computing for Researchers: Laptop

- **Regular PC / Local Laptop / Workstation**
  ↪ **Native OS** (Windows, Linux, Mac etc.)

Computing for Researchers: Laptop

- **Regular PC / Local Laptop / Workstation**
  - ↪ **Native OS** (Windows, Linux, Mac etc.)
- Virtualized OS (**VM**) through an **hypervisor**
  - ↪ *Hypervisor*: core virtualization engine / environment
    - ✓ Ex: Xen, VMWare ESXi, KVM, VirtualBox
    - ✓ Non-negligible Performance loss: $\geq 20\%$

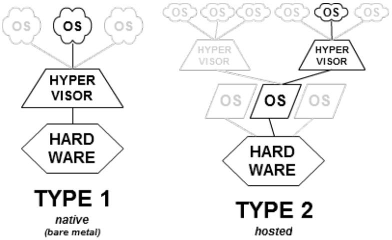

Computing for Researchers: Laptop

- **Regular PC / Local Laptop / Workstation**
  - ↪ **Native OS** (Windows, Linux, Mac etc.)
- Virtualized OS (**VM**) through an **hypervisor**
  - ↪ *Hypervisor*: core virtualization engine / environment
    - ✓ Ex: Xen, VMWare ESXi, KVM, VirtualBox
    - ✓ Non-negligible Performance loss: $\geq$ 20%
- **Container**-based Virtualization
  - ↪ similar to VMs ...
    - ✓ **yet** containers **share** the system kernel of the host with others
    - ✓ Ex: Docker, Singularity, Shifter

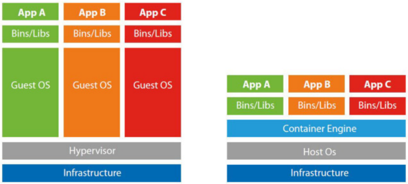

Computing for Researchers: Cloud

- **Cloud Computing**
  - ↪ access to shared *(generally virtualized)* resources
  - ↪ pay-per-use approach
  - ↪ **Infrastructure** as a Service (**I**aaS)

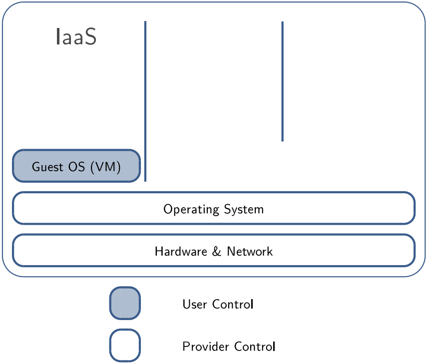

Computing for Researchers: Cloud

- **Cloud Computing**
  - ↪ access to shared *(generally virtualized)* resources
  - ↪ pay-per-use approach
  - ↪ **Platform** as a Service (**P**aaS)

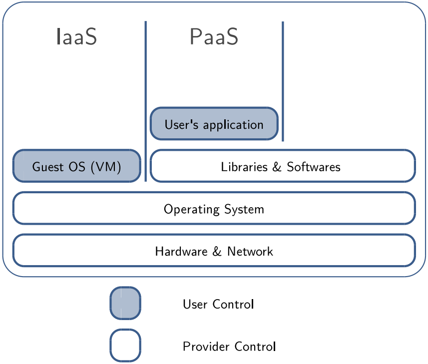

Computing for Researchers: Cloud

- **Cloud Computing**
  - ↪ access to shared *(generally virtualized)* resources
  - ↪ pay-per-use approach
  - ↪ **Software** as a Service (**S**aaS)

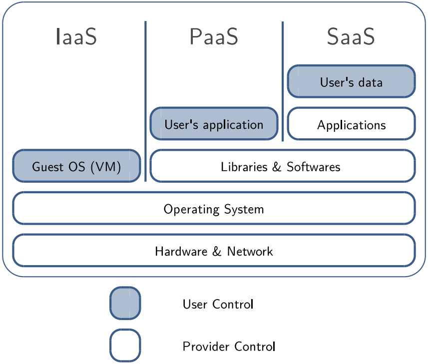

Computing for Researchers: Cloud

- **Cloud Computing**
  - ↪ access to shared *(generally virtualized)* resources
  - ↪ pay-per-use approach
  - ↪ **XXX** as a Service (<**X**>aaS)

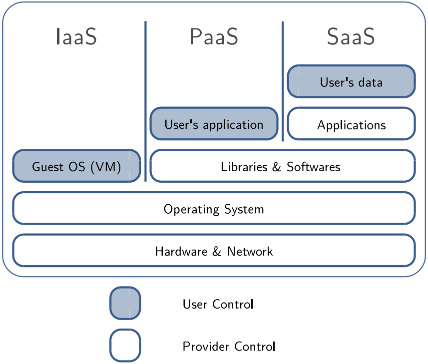

Computing for Researchers: HPC

- High Performance Computing (HPC) platforms
  - ↪ For **Speedup**, **Scalability** and **Faster Time to Solution**

Computing for Researchers: HPC

- High Performance Computing (HPC) platforms
  ↪ For **Speedup**, **Scalability** and **Faster Time to Solution**

YET...

$$\textbf{PC} \neq \textbf{Cloud} \neq \textbf{HPC}$$

Computing for Researchers: HPC

- High Performance Computing (HPC) platforms
  - ↪ For **Speedup**, **Scalability** and **Faster Time to Solution**

YET...

PC $\neq$ Cloud $\neq$ HPC

- HPC $\simeq$ Formula 1
  - ↪ relies on ultra efficient hardware / interconnect (IB EDR. . .)
  - ↪ . . . when Cloud has to stay standard ([10] GbE etc. . .)
- **Does not mean the 3 approaches cannot work together**

Different HPC Needs per Domains

Material Science & Engineering

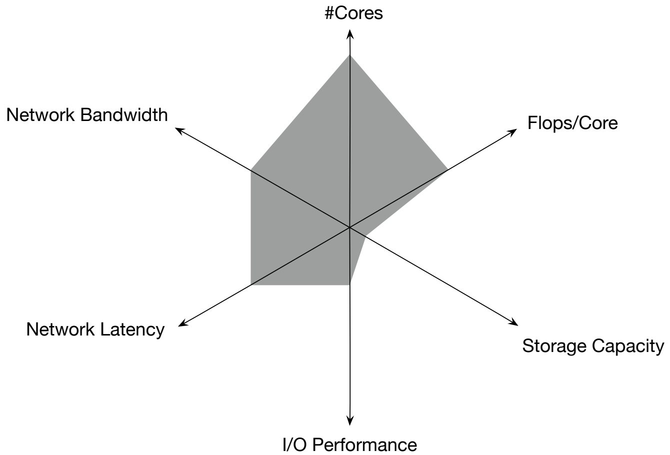

Different HPC Needs per Domains

Biomedical Industry / Life Sciences

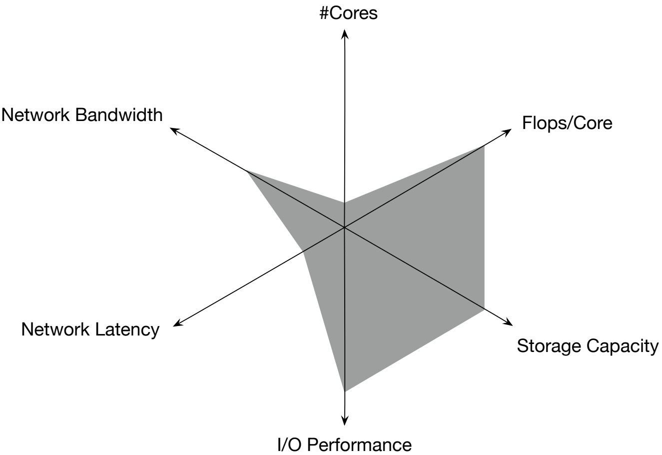

Different HPC Needs per Domains

Deep Learning / Cognitive Computing

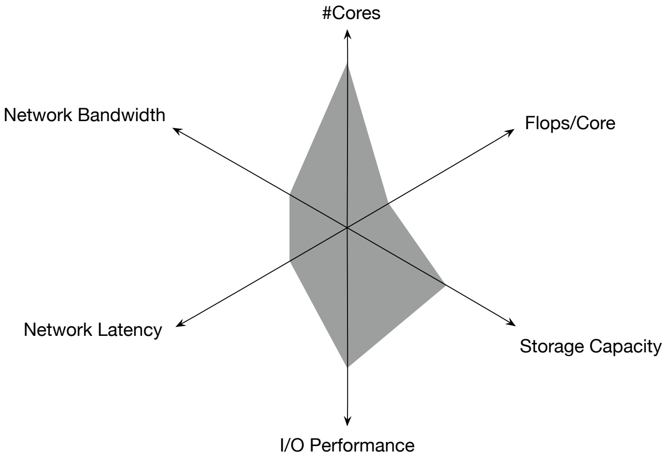

Different HPC Needs per Domains

IoT, FinTech

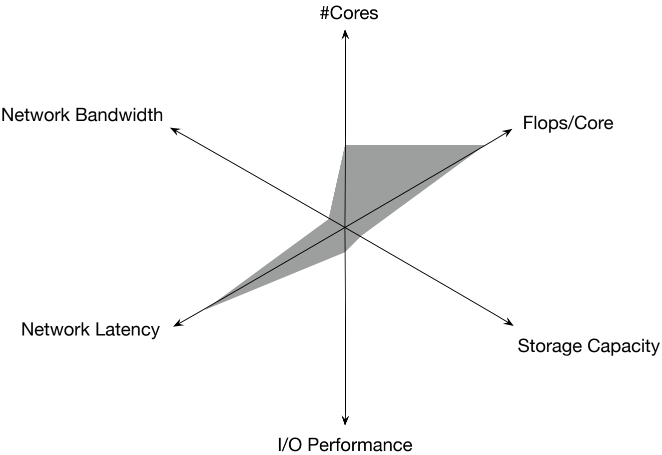

Different HPC Needs per Domains

ALL Research Computing Domains

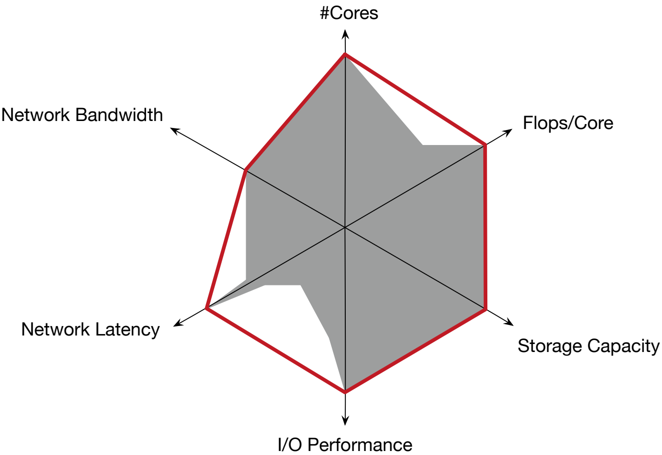

Summary

HPC Components: [GP]CPU

CPU

- Always multi-core
- Ex: Intel Core i7-7700K (Jan 2017) $R_{peak} \simeq$ 268.8 GFlops (DP)
  - ↪ 4 cores @ 4.2GHz (14nm, 91W, 1.75 billion transistors)
  - ↪ + integrated graphics (24 EUs) $R_{peak} \simeq$ +441.6 GFlops

GPU / GPGPU

- Always multi-core, optimized for vector processing
- Ex: Nvidia Tesla V100 (Jun 2017) $R_{peak} \simeq$ 7 TFlops (DP)
  - ↪ 5120 cores @ 1.3GHz (12nm, 250W, 21 billion transistors)
  - ↪ focus on Deep Learning workloads $R_{peak} \simeq$ 112 TFLOPS (HP)

$\simeq$ **100 Gflops** for 130$ (CPU), 214$? (GPU)

HPC Components: Local Memory

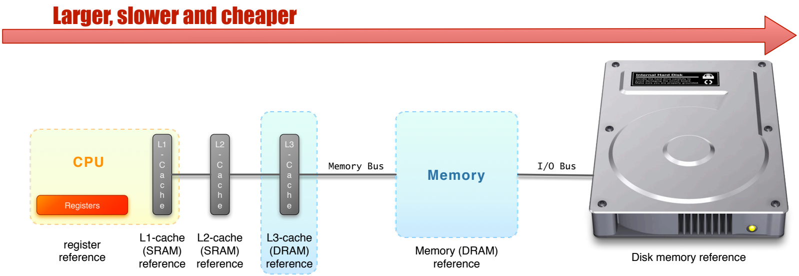

Level:	1	2	3	4
Size:	500 bytes	64 KB to 8 MB	1 GB	1 TB
Speed:	sub ns	1-2 cycles 10 cycles 20 cycles	hundreds cycles	ten of thousands cycles

- SSD (SATA3) R/W: 550 MB/s; 100000 IOPS **450 €/TB**
- HDD (SATA3 @ 7,2 krpm) R/W: 227 MB/s; 85 IOPS **54 €/TB**

HPC Components: Interconnect

- **latency**: time to send a minimal (0 byte) message from A to B
- **bandwidth**: max amount of data communicated per unit of time

Technology	Effective Bandwidth		Latency
Gigabit Ethernet	1 Gb/s	125 MB/s	40μs to 300μs
10 Gigabit Ethernet	10 Gb/s	1.25 GB/s	4μs to 5μs
Infiniband QDR	40 Gb/s	5 GB/s	1.29μs to 2.6μs
Infiniband EDR	100 Gb/s	12.5 GB/s	0.61μs to 1.3μs
100 Gigabit Ethernet	100 Gb/s	1.25 GB/s	30μs
Intel Omnipath	100 Gb/s	12.5 GB/s	0.9μs

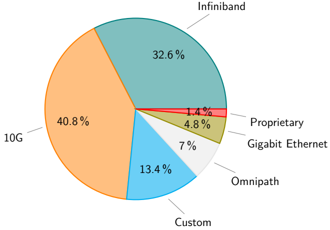

[Source : www.top500.org, Nov. 2017]

HPC Components: Interconnect

- **latency**: time to send a minimal (0 byte) message from A to B
- **bandwidth**: max amount of data communicated per unit of time

Technology	Effective Bandwidth		Latency
Gigabit Ethernet	1 Gb/s	125 MB/s	40μs to 300μs
10 Gigabit Ethernet	10 Gb/s	1.25 GB/s	4μs to 5μs
Infiniband QDR	40 Gb/s	5 GB/s	1.29μs to 2.6μs
Infiniband EDR	100 Gb/s	12.5 GB/s	0.61μs to 1.3μs
100 Gigabit Ethernet	100 Gb/s	1.25 GB/s	30μs
Intel Omnipath	100 Gb/s	12.5 GB/s	0.9μs

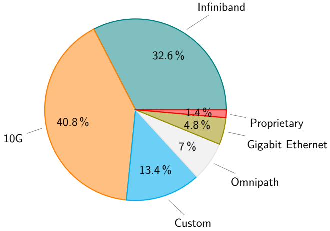

[Source : www.top500.org, Nov. 2017]

Network Topologies

- **Direct** vs. **Indirect** interconnect
  - ↪ *direct*: each network node attaches to at least one compute node
  - ↪ *indirect*: compute nodes attached at the edge of the network only
    - ✓ many routers only connect to other routers.

Network Topologies

- **Direct** vs. **Indirect** interconnect
  - ↪ *direct*: each network node attaches to at least one compute node
  - ↪ *indirect*: compute nodes attached at the edge of the network only
    - ✓ many routers only connect to other routers.

Main HPC Topologies

- **CLOS Network / Fat-Trees** [Indirect]
  - ↪ can be fully non-blocking (1:1) or blocking (x:1)
  - ↪ typically enables **best performance**
    - ✓ Non blocking bandwidth, lowest network latency

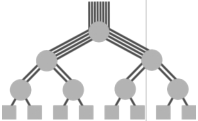

Network Topologies

- **Direct** vs. **Indirect** interconnect
  - ↪ *direct*: each network node attaches to at least one compute node
  - ↪ *indirect*: compute nodes attached at the edge of the network only
    - ✓ many routers only connect to other routers.

Main HPC Topologies

- **CLOS Network / Fat-Trees** [Indirect]
  - ↪ can be fully non-blocking (1:1) or blocking (x:1)
  - ↪ typically enables **best performance**
    - ✓ Non blocking bandwidth, lowest network latency

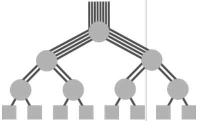

- **Mesh or 3D-torus** [Direct]
  - ↪ Blocking network, cost-effective for systems at scale
  - ↪ Great performance solutions for applications with locality
  - ↪ Simple expansion for future growth

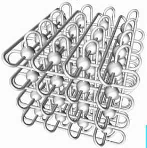

HPC Components: Operating System

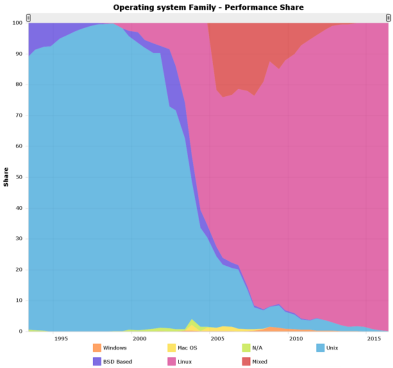

- Exclusively Linux-based (**really** 100%)
- Reasons:
  - ↪ stability
  - ↪ prone to devels

[Source : www.top500.org, Nov 2017]

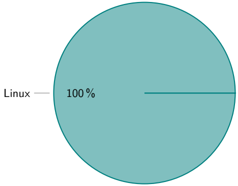

HPC Components: Software Stack

- **Remote connection to the platform** SSH
- **Identity Management / SSO**: LDAP, Kerberos, IPA...
- **Resource management**: job/batch scheduler
  ↪ SLURM, OAR, PBS, MOAB/Torque...
- (Automatic) **Node Deployment**:
  ↪ FAI, Kickstart, Puppet, Chef, Ansible, Kadeploy...
- (Automatic) **User Software Management**:
  ↪ Easybuild, Environment Modules, LMod
- **Platform Monitoring**:
  ↪ Nagios, Icinga, Ganglia, Foreman, Cacti, Alerta...

[Big]Data Management: Disk Encl.

- $\simeq$ **120 K€** - enclosure - 48-60 disks (4U)
  - ↪ incl. redundant (**i.e.** 2) RAID controllers (master/slave)

[Big]Data Management: FS Summary

- **File System (FS)**: Logical manner to *store, organize & access* data
  - ↪ (local) **Disk FS** : `FAT32`, `NTFS`, `HFS+`, `ext4`, `{x,z,btr}fs`...
  - ↪ **Networked FS**: `NFS`, `CIFS/SMB`, `AFP`
  - ↪ **Parallel/Distributed FS**: `SpectrumScale/GPFS`, `Lustre`
    - ✓ typical FS for HPC / HTC (High Throughput Computing)

[Big]Data Management: FS Summary

- **File System (FS)**: Logical manner to *store, organize & access* data
  - ↪ (local) **Disk FS** : `FAT32`, `NTFS`, `HFS+`, `ext4`, `{x,z,btr}fs`...
  - ↪ **Networked FS**: `NFS`, `CIFS/SMB`, `AFP`
  - ↪ **Parallel/Distributed FS**: `SpectrumScale/GPFS`, `Lustre`
    - ✓ typical FS for HPC / HTC (High Throughput Computing)

Main Characteristic of Parallel/Distributed File Systems

Capacity and Performance increase with #servers

[Big]Data Management: FS Summary

- **File System (FS)**: Logical manner to *store, organize & access* data
  - ↪ (local) **Disk FS** : FAT32, NTFS, HFS+, ext4, {x,z,btr}fs...
  - ↪ **Networked FS**: NFS, CIFS/SMB, AFP
  - ↪ **Parallel/Distributed FS**: SpectrumScale/GPFS, Lustre
    - ✓ typical FS for HPC / HTC (High Throughput Computing)

Main Characteristic of Parallel/Distributed File Systems

Capacity and Performance increase with #servers

Name	Type	Read* [GB/s]	Write* [GB/s]
ext4	Disk FS	0.426	0.212
nfs	Networked FS	0.381	0.090
gpfs (iris)	Parallel/Distributed FS	**11.25**	**9,46**
lustre (iris)	Parallel/Distributed FS	**12.88**	**10,07**
gpfs (gaia)	Parallel/Distributed FS	7.74	6.524
lustre (gaia)	Parallel/Distributed FS	4.5	2.956

* maximum **random** read/write, per IOZone or IOR measures, using concurrent nodes for networked FS.

HPC Components: Data Center

Definition (Data Center)

- Facility to house computer systems and associated components
  $\hookrightarrow$ Basic storage component: **rack** (height: 42 RU)

HPC Components: Data Center

Definition (Data Center)

- Facility to house computer systems and associated components
  - ↪ Basic storage component: **rack** (height: 42 RU)

Challenges: Power (UPS, battery), Cooling, Fire protection, Security

- Power/Heat dissipation per rack:
  - ↪ HPC **computing** racks: **30-120 kW**
  - ↪ **Storage** racks: **15 kW**
  - ↪ **Interconnect** racks: **5 kW**
- Various **Cooling** Technology
  - ↪ Airflow
  - ↪ Direct-Liquid Cooling, Immersion...

Power Usage Effectiveness

$$PUE = \frac{\textit{Total facility power}}{\textit{IT equipment power}}$$

HPC Components: Summary

Running an HPC Facility involves...

- A **data center** / server room carefully designed
- Many **computing** elements: CPU, GPGPU, Accelerators
- **Fast interconnect** elements
  - ↪ high *bandwidth* and low *latency*
- [Big]-Data **storage** elements: HDD/SDD, disk enclosure,
  - ↪ disks are virtually aggregated by RAID/LUNs/FS
  - ↪ parallel and distributed FS
- A flexible software stack
- Automated management everywhere

Above all: **expert** HPC/IT specialists !

Summary

Univ. of Luxembourg & HPC

- *With regards to HPC*, Univ. of Luxembourg offers:
  - ↪ **People**
    - ✓ Domain experts
    - ✓ Computational and data scientists
    - ✓ Specialists in parallel algorithmics
  - ↪ **Services**
    - ✓ HPC clusters and management team
    - ✓ IT team (SIU)
    - ✓ Infrastructure team in collab. w. Fonds Belval
  - ↪ **Infrastructure**
    - ✓ Data center and a set high-end clusters
  - ↪ **Education & Training**

High Performance Computing @ UL

- **Started in 2007**
  - ↪ under resp. of Prof P. Bouvry & Dr. S. Varrette
  - ↪ expert UL HPC team
    - ✓ *S. Varrette, V. Plugaru, S. Peter, H. Cartiaux, C. Parisot*. . . and multiple domain experts per RU
  - ↪ Largest HPC facility in Luxembourg w. GoodYear

https://hpc.uni.lu

High Performance Computing @ UL

- **Started in 2007**
  - ↪ under resp. of Prof P. Bouvry & Dr. S. Varrette
  - ↪ expert UL HPC team
    - ✓ *S. Varrette, V. Plugaru, S. Peter, H. Cartiaux, C. Parisot*... and multiple domain experts per RU
  - ↪ Largest HPC facility in Luxembourg w. GoodYear

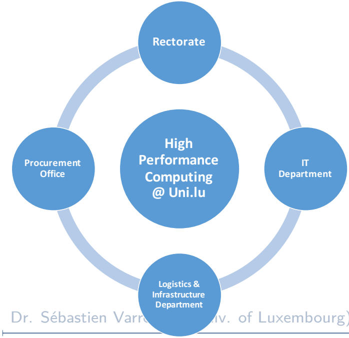

High Performance Computing @ UL

- **Started in 2007**
  - ↪ under resp. of Prof P. Bouvry & Dr. S. Varrette
  - ↪ expert UL HPC team
    - ✓ *S. Varrette, V. Plugaru, S. Peter, H. Cartiaux, C. Parisot*... and multiple domain experts per RU
  - ↪ Largest HPC facility in Luxembourg w. GoodYear

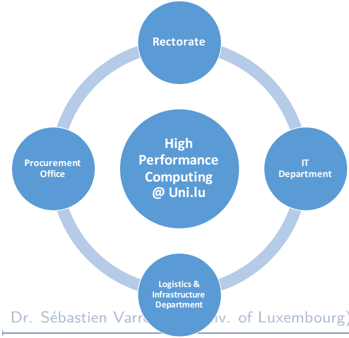

HPC/Computing Capacity

1029.342 TFlops

(incl. 612.62 GPU TFlops)

HTC/Storage Capacity

9852.4 TB storage

UL HPC Computing capacity

5 clusters / 2 sites

1029.342 TFlops

(incl. 612.62 GPU TFlops)

684 nodes

11084 CPU cores

(+ 489344 GPU cores)

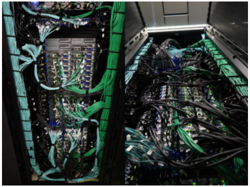

- IB interconnect
- Fat tree topo. in general

INFINIBAND TRADE ASSOCIATION

UNIVERSITÉ DU LUXEMBOURG

UL HPC Storage capacity

9852.4 TB (incl. 1020TB for Backup)

2425 disks

- 4 distributed/parallel FS
  - ↪ GPFS : 3244 TB
  - ↪ Lustre: 1940 TB
  - ↪ OneFS: 3188 TB...

UL HPC Beneficiaries

23 computational domains accelerated on UL HPC

- for the UL Faculties, Research Units and Interdisciplinary Centres
  - ↪ incl. LCSB, SnT... and now C2DH thematics
  - ↪ UL **strategic research priorities**
    - ✓ computational sciences, finance (fintech)
    - ✓ systems biomedicine, security, reliability and trust

- UL HPC feat. special systems targeting specific workloads:
  - ↪ **Machine Learning & AI**: GPU accelerators
    - ✓ 10 Tesla K40 + 16 Tesla K80 + 24 Tesla M20*: **76 GPU Tflops**
    - ✓ **Q4 2018**: 18*4 V100 (part of RFP 180027): **> 365 GPU Tflops**
  - ↪ **BigData analytics & data driven science**: large memory systems
    - ✓ Large SMP systems with 1, 2, 3 & 4 TB RAM
  - ↪ **Scale-out workloads**: 90 HP Moonshot servers + 96 viridis ARM-based systems

Accelerating UL Research

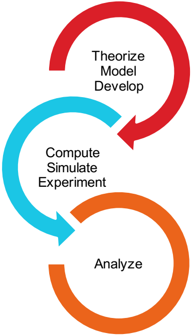

- **198 software packages** available for researchers
  - ↪ via Environment modules/LMod from Easybuild

Domain	**Software**
Compiler Toolchains	(2018a) FOSS, Intel, PGI
MPI suits	OpenMPI, Intel MPI, MVAPICH2
Machine Learning	PyTorch, TensorFlow, Keras, Apache Spark. . .
Math & Optimization	Matlab, Mathematica, R, CPLEX. . .
Physics & Chemistry	GROMACS, QuantumESPRESSO, ABINIT, NAMD, VASP. . .
Bioinformatics	SAMtools, BLAST+, ABySS, mpiBLAST, TopHat, Bowtie2. . .
Computer aided engineering	ANSYS, ABAQUS, OpenFOAM. . .
General purpose	Allinea/ARM Forge & Perf Reports, Python, Go, Rust. . .
Container systems	Singularity
Visualisation	ParaView, OpenCV, XCS portal
. . .	

`https://hpc.uni.lu/users/software/`

Computing nodes Management

Node deployment by FAI/Bright Manager

- Boot via network card (PXE)
  - ↪ ensure a running diskless Linux OS

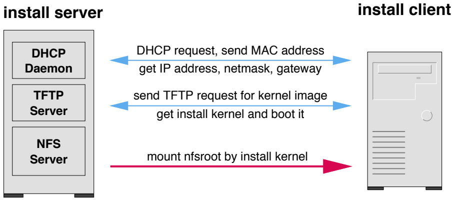

Computing nodes Management

Node deployment by FAI/Bright Manager

- Boot via network card (PXE)
  - ↪ ensure a running diskless Linux OS
- Get configuration data (NFS/other)

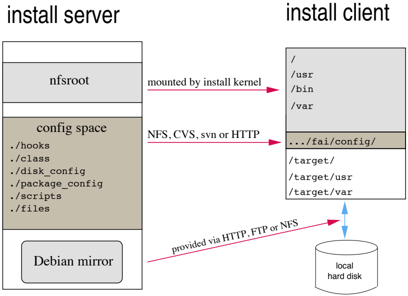

Computing nodes Management

Node deployment by FAI/Bright Manager

- Boot via network card (PXE)
  - ↪ ensure a running diskless Linux OS
- Get configuration data (NFS/other)
- Run the installation
  - ↪ partition local hard disks and create filesystems
  - ↪ install software using apt-get command
  - ↪ configure OS and additional software
  - ↪ save log files to install server, then reboot new system

Computing nodes Management

Node deployment by FAI/Bright Manager

- Boot via network card (PXE)
  - ↪ ensure a running diskless Linux OS
- Get configuration data (NFS/other)
- Run the installation
  - ↪ partition local hard disks and create filesystems
  - ↪ install software using apt-get command
  - ↪ configure OS and additional software
  - ↪ save log files to install server, then reboot new system

Average reinstallation time: $\simeq$ 500s

IT Serv[er|ice] Management: Puppet

Server/Service configuration by Puppet

http://puppetlabs.com

- **IT Automation** for configuration management
  - ↪ idempotent, agent/master OR stand-alone
  - ↪ cross-platform through Puppet Resource Abstraction Layer (RAL)
  - ↪ git-based workflow with r10k (**role & profiles** workflow)
  - ↪ PKI-based security (X.509)
- DevOps tool of choice for configuration management
  - ↪ Reusable modules
  - ↪ per-environment hierarchy lookup with hiera

https://forge.puppet.com/

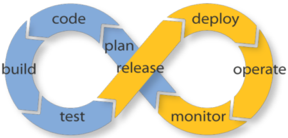

Endless Possibilities: DevOps can create an infinite loop of release and feedback for all your code and deployment targets.

IT Serv[er|ice] Management: Puppet

Server/Service configuration by Puppet

http://puppetlabs.com

- **IT Automation** for configuration management
  - ↪ idempotent, agent/master OR stand-alone
  - ↪ cross-platform through Puppet Resource Abstraction Layer (RAL)
  - ↪ git-based workflow with r10k (**role & profiles** workflow)
  - ↪ PKI-based security (X.509)
- DevOps tool of choice for configuration management
  - ↪ Reusable modules https://forge.puppet.com/
  - ↪ per-environment hierarchy lookup with hiera

Average server installation/configuration time: $\simeq$ 3-6 min

Summary

HPC International State of Affairs

Global race toward Exascale Technology

IDC-Projected Exascale Investment Levels (In Addition to System Purchases)

U.S.
- $1 to $2 billion a year in R&D (including NRE)
- Investments by both governments & vendors
- Plans are to purchase multiple exascale systems

EU
- About 5 billion euros in total
- Investments in multiple exascale and pre-exascale systems
- Investments mostly by country governments with a little from the EU

China
- Over $1billion a year in R&D
- Investments by both governments & vendors
- Plans are to purchases multiple exascale systems each year
- Already investing in 3 pre-exascale systems by 2017/18

Japan
- Planned investment of just over $1billion* (over 5 years) for both the R&D and purchase of 1 exascale system
- To be followed by a number of smaller systems ~$100M to $150M each
- Creating a new processor and a new software environment

Note that this includes both the system and R&D

HPC International State of Affairs

Global race toward Exascale Technology

IDC-Projected Exascale Dates and Suppliers

U.S.
- Sustained ES: 2023
- Peak ES: 2021
- Vendors: U.S.
- Processors: U.S.
- Initiatives: NSCI/ECP
- Cost: $300-500M per system, plus heavy R&D investments

EU
- Sustained ES: 2023-24
- Peak ES: 2021
- Vendors: U.S., Europe
- Processors: U.S., ARM
- Initiatives: PRACE, ETP4HPC
- Cost: $300-$350 per system, plus heavy R&D investments

China
- Sustained ES: 2023
- Peak ES: ~~2020~~ 2019…
- Vendors: Chinese
- Processors: Chinese (plus ~~U.S.~~?)
- 13th 5-Year Plan
- Cost: $350-500M per system, plus heavy R&D

Japan
- Sustained ES: 2023-24
- Peak ES: Not planned
- Vendors: Japanese
- Processors: Japanese
- Cost: $600-850M, this includes both 1 system and the R&D costs…will also do many smaller size systems

Exascale Feasibility

We Can Build an Exascale System Today?

Connect together 10 Sunway TaihuLight systems

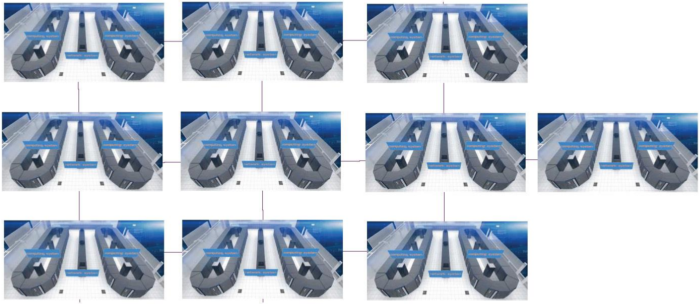

Require 150 MW of power, programming for 100 M threads, and $2.7B price tag

22

European HPC strategy

- EU HPC strategy initiated in 2012
  ↪ implementation within H2020 program

European HPC strategy


IMPORTANT PROJECT OF COMMON EUROPEAN INTEREST (IPCEI)

ON HIGH PERFORMANCE COMPUTING AND BIG DATA ENABLED APPLICATIONS (IPCEI-HPC-BDA)

European Strategic Positioning Paper

Luxembourg, France, Italy (& Spain)

November 2015

- EU HPC strategy initiated in 2012
  - ↪ implementation within H2020 program
- More recently:
  - ↪ IPCEI on HPC and Big Data (BD) Applications (Nov. 2015)
    - ✓ Luxembourg (leader), France, Italy & Spain
    - ✓ Testbed around Personalized Medicine, Smart Space, Industry 4.0, Smart Manufacturing, New Materials, FinTech, Smart City...

European HPC strategy

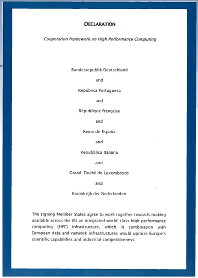
DECLARATION

Cooperation framework on High Performance Computing

Bundesrepublik Deutschland

and

República Portuguesa

and

République française

and

Reino de España

and

Repubblica Italiana

and

Grand-Duché de Luxembourg

and

Koninkrijk der Nederlanden

The signing Member States agree to work together towards making available across the EU an integrated world-class high performance computing (HPC) infrastructure, which in combination with European data and network infrastructures would upraise Europe's scientific capabilities and industrial competitiveness.

- EU HPC strategy initiated in 2012
  - ↪ implementation within H2020 program
- More recently:
  - ↪ IPCEI on HPC and Big Data (BD) Applications
    - ✓ Luxembourg (leader), France, Italy & Spain
    - ✓ Testbed around Personalized Medicine, Smart Space, Industry 4.0, Smart Manufacturing, New Materials, FinTech, Smart City...
- **Latest advances**:
  - ↪ EU Member States sign EuroHPC (Mar. 2017)
    - ✓ common effort to create/grow European supercomputing ecosystem
    - ✓ Federation of national/regional HPC centers (see also PRACE2)
  - ↪ EU Objective with EuroHPC:
    - ✓ EuroHPC JU effectively operational starting **Jan 1st, 2019**
    - ✓ 2-3 **Pre-exascale** systems 2020, **2 exascale** systems by 2022

EU HPC Strategy Implementation

- **European Technology Platform (ETP) for HPC**
  - ↪ Industry-led forum feat. HPC stakeholders
  - ↪ Providing EU framework to define HPC research priorities/actions
    - ✓ UL (P. Bouvry, S. Varrette, V.Plugaru) part of ETP4HPC (2016-)
    - ✓ See Strategic Research Agenda, 2017 European HPC Handbook. . .

EU HPC Strategy Implementation

- **European Technology Platform (ETP) for HPC**
  - ↪ Industry-led forum feat. HPC stakeholders
  - ↪ Providing EU framework to define HPC research priorities/actions
    - ✓ UL (P. Bouvry, S. Varrette, V.Plugaru) part of ETP4HPC (2016-)
    - ✓ See Strategic Research Agenda, 2017 European HPC Handbook. . .

- **PRACE** - Partnership for Advanced Computing in Europe
  - ↪ Non-profit association, 25 member countries, now entering PRACE2
  - ↪ Providing access to **Five EU Tier-0** compute & data resources
  - ↪ (Oct. 2017) **Luxembourg 25th country to join PRACE**
    - ✓ Official Delegate/Advisor (P. Bouvry/S. Varrette) from UL

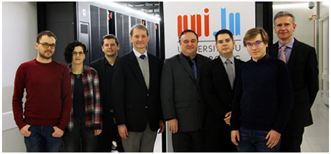

EU HPC Strategy Implementation

- **European High-Performance Computing Joint Undertaking**
  - ↪ EuroHPC JU effectively operational starting **Jan 1st, 2019**
    - ✓ administrative management from Luxembourg
  - ↪ Public and private members
    - ✓ EC, 14 MS, representatives from supercomputing/BD stakeholders
    - ✓ Governing Board (public members)
    - ✓ Industrial & Scientific Advisory Board (private members)
  - ↪ EU Objective with EuroHPC:
    - ✓ 2-3 **Pre-exascale** systems (2020), **2 exascale** systems (2022)
    - ✓ Pending decision on hosting countries

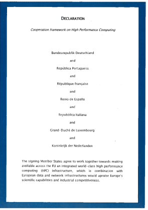

DECLARATION

Cooperation framework on High Performance Computing

Bundesrepublik Deutschland

and

República Portuguesa

and

République française

and

Reino de España

and

Repubblica Italiana

and

Grand-Duché de Luxembourg

and

Koninkrijk der Nederlanden

The signing Member States agree to work together towards making available across the EU an integrated world-class high performance computing (HPC) infrastructure, which in combination with European data and network infrastructures would upraise Europe's scientific capabilities and industrial competitiveness.

EU HPC Strategy Implementation

- **European High-Performance Computing Joint Undertaking**
  - ↪ EuroHPC JU effectively operational starting **Jan 1st, 2019**
    - ✓ administrative management from Luxembourg
  - ↪ Public and private members
    - ✓ EC, 14 MS, representatives from supercomputing/BD stakeholders
    - ✓ Governing Board (public members)
    - ✓ Industrial & Scientific Advisory Board (private members)
  - ↪ EU Objective with EuroHPC:
    - ✓ 2-3 **Pre-exascale** systems (2020), **2 exascale** systems (2022)
    - ✓ Pending decision on hosting countries

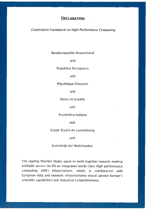

DECLARATION

Cooperation framework on High Performance Computing

Bundesrepublik Deutschland

and

República Portuguesa

and

République française

and

Reino de España

and

Repubblica Italiana

and

Grand-Duché de Luxembourg

and

Koninkrijk der Nederlanden

The signing Member States agree to work together towards making available across the EU an integrated world-class high performance computing (HPC) infrastructure, which in combination with European data and network infrastructures would upraise Europe's scientific capabilities and industrial competitiveness.

EuroHPC Budget: 2 × 486 M€

EU HPC Strategy Implementation

- **European High-Performance Computing Joint Undertaking**
  - ↪ EuroHPC JU effectively operational starting **Jan 1st, 2019**
    - ✓ administrative management from Luxembourg
  - ↪ Public and private members
    - ✓ EC, 14 MS, representatives from supercomputing/BD stakeholders
    - ✓ Governing Board (public members)
    - ✓ Industrial & Scientific Advisory Board (private members)
  - ↪ EU Objective with EuroHPC:
    - ✓ 2-3 **Pre-exascale** systems (2020), **2 exascale** systems (2022)
    - ✓ Pending decision on hosting countries

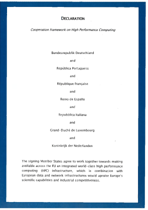
DECLARATION

Cooperation framework on High Performance Computing

Bundesrepublik Deutschland

and

República Portuguesa

and

République française

and

Reino de España

and

Repubblica Italiana

and

Grand-Duché de Luxembourg

and

Koninkrijk der Nederlanden

The signing Member States agree to work together towards making available across the EU an integrated world-class high performance computing (HPC) infrastructure, which in combination with European data and network infrastructures would upraise Europe's scientific capabilities and industrial competitiveness.

EuroHPC Budget: 2 × 486 M€

- **European Processor Initiative (EPI)**
  - ↪ Initial plan vs current plan. . .
  - ↪ **120 M€** via Framework Partnership Agreement (FPA)

Summary

Summary

New Trends in HPC

- **Continued scaling** of scientific, industrial & financial applications
  - ↪ . . . well beyond Exascale
- New trends changing the landscape for HPC
  - ↪ Emergence of **Big Data analytics**
  - ↪ Emergence of (**Hyperscale**) **Cloud Computing**
  - ↪ **Data intensive Internet of Things (IoT)** applications
  - ↪ **Deep learning & cognitive computing** paradigms

[Source : EuroLab-4-HPC]


This study was carried out for RIKEN by

IDC Analyze the Future

Special Study

Analysis of the Characteristics and Development Trends of the Next-Generation of Supercomputers in Foreign Countries

Earl C. Joseph, Ph.D.
Steve Conway

Robert Sorensen
Kevin Monroe

[Source : IDC RIKEN report, 2016]

Toward Modular Computing

- Aiming at **scalable, flexible HPC infrastructures**
  - ↪ *Primary* processing on CPUs and accelerators
    - ✓ **HPC & Extreme Scale Booster** modules
  - ↪ *Specialized modules* for:
    - ✓ **HTC & I/O intensive** workloads;
    - ✓ **[Big] Data Analytics & AI**

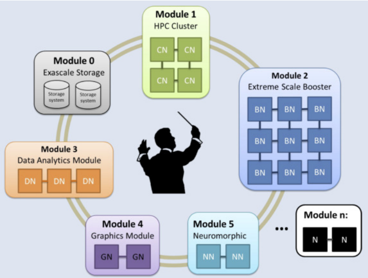

[Source : "Towards Modular Supercomputing: The DEEP and DEEP-ER projects", 2016]

Summary

Uni.lu CDC (Centre de Calcul)

- **Toward Energy-Efficient HPC enabling DLC**
  - ↪ 2x500 m^2 deployed since 2015, one floor for HPC developments

Location	Cooling	Usage	Max Capacity [kW]
CDC S-02-001	Airflow	*Future extension*	280 kW (120 m2)
CDC S-02-002	Airflow	Storage / Traditional HPC /Cloud/FPGA	280 kW (88 m2)
CDC S-02-003	**DLC**	**High Density/Energy efficient HPC**	**1050 kW (90 m2)**
CDC S-02-004	**DLC**	**High Density/Energy efficient HPC**	**1050 kW (92 m2)**
CDC S-02-005	Airflow	Storage / Traditional HPC (iris cluster)	300 kW (128 m2)

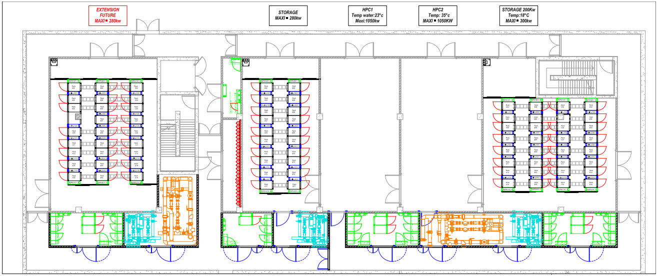

National HPC-BD Competence Center

- Built by ministerial, academic, industrial stakeholders
  - ↪ **Inspired by national research computing centers**

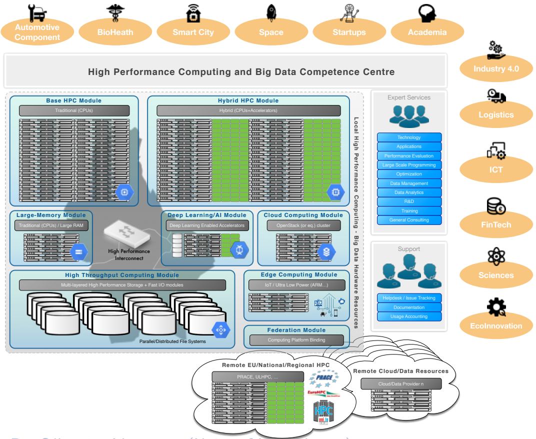

- Comprehensive centre:
  - ↪ HPC
  - ↪ data infrastructure
  - ↪ Techn. Expertise
  - ↪ Domain knowledge
- *More than just computing services*
- **Inspiration**:
  - ↪ **EU**: JSC, TGCC...
  - ↪ **US**: OSC, SDSC, TACC, LLNL...
  - ↪ **Singapore**: A*STAR

Summary

Quantum Computing

- Based on the **laws of quantum mechanics**
  - ↪ behaviour of particles at [sub-]atomic level
- Electrons (or any other subatomic particle) has super powers:
  - ↪ **superposition**: be a wave and a particle, at the same time
  - ↪ **quantum tunnelling**: bypass any barriers *i.e* move through walls
  - ↪ **entanglement**: power to be psychic like having a twin
    - ✓ where if one is affected then simultaneously so is the other

Quantum Computing

- Based on the **laws of quantum mechanics**
  - ↪ behaviour of particles at [sub-]atomic level
- Electrons (or any other subatomic particle) has super powers:
  - ↪ **superposition**: be a wave and a particle, at the same time
  - ↪ **quantum tunnelling**: bypass any barriers *i.e* move through walls
  - ↪ **entanglement**: power to be psychic like having a twin
    - ✓ where if one is affected then simultaneously so is the other

Why so much buzz???

- Enables to compute huge amounts of data simultaneously
  - ↪ ... and solve complex problems such as optimization
- Ex: general number field sieve (GNFS) vs. **Shor's** Algorithm
  - ↪ shows (in principle) capability of factoring very large numbers
    - ✓ ... in polynomial time

Data Representation: Qubits

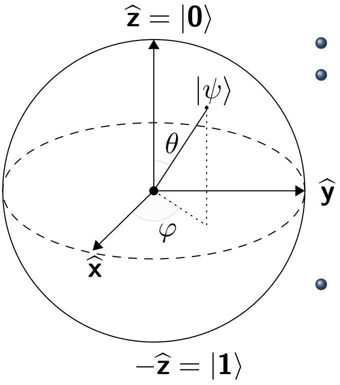

- **Classical** bit: state 0 **or** 1
- **Quantum bit** (Qubit):
  - ↪ a single atom in one of two states $|0\rangle$ or $|1\rangle$
  - ↪ single qubit can be forced into a **superposition** of two states

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$$

- Qubit in superposition is in **both** states
  - ↪ ... at the **same** time
  - ↪ in general, *n* **qubits register** can represent the numbers in $[0, 2^n - 1]$ **simultaneously**

Data Representation: Qubits

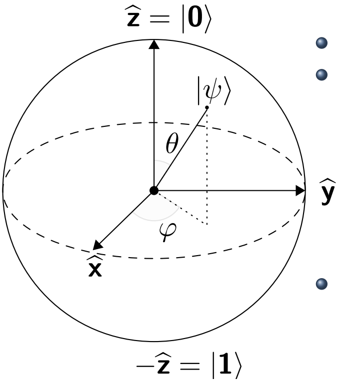

- **Classical** bit: state 0 **or** 1
- **Quantum bit** (Qubit):
  - ↪ a single atom in one of two states $|0\rangle$ or $|1\rangle$
  - ↪ single qubit can be forced into a **superposition** of two states

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$$

- Qubit in superposition is in **both** states
  - ↪ ... at the **same** time
  - ↪ in general, *n* **qubits register** can represent the numbers in $[0, 2^n - 1]$ **simultaneously**

Reading value ⇒ random collapse of superposition to one state

Quantum Algorithm

- Relies on **quantum gates** (reversible logic)
  - ↪ Ex: Hadamard gate $\sqrt{NOT}$, controlled NOT (CN) $\simeq XOR$...

$\begin{bmatrix}1\\0\\0\\\vdots\\0\end{bmatrix}$ State Preparation $\begin{bmatrix}\alpha_0\\\alpha_1\\\alpha_2\\\vdots\\\alpha_n\end{bmatrix}$ Quantum Algorithm — Measuring Qbit 0 $\begin{bmatrix}0\\\beta'_1\\0\\\vdots\\\beta'_n\end{bmatrix}$ Measuring remaining Qbits $\begin{bmatrix}0\\1\\0\\\vdots\\0\end{bmatrix}$

Time →

Number of Qbits	Number of superposed states
1	2
2	4
n	2^n

TOP 500 place	Amount of RAM	Number of Qbits
1	2.8 PB	47
2	1.3 PB	46

Courtesy of Dr. Jean-Noël Quintin, Quantum Computing Lab, Atos/Bull

Quantum Algorithm

1. Activate qubits to reach a superpositions of all possible states
2. Encode the optimization problem
   - ↪ by applying a phase on each superposition state
3. Use methods of interference to cancel or add phases to optimize for the correct answer
   - ↪ and shrink the wrong answers (like noise canceling in headphones)

Quantum Algorithm

1. Activate qubits to reach a superpositions of all possible states
2. Encode the optimization problem
   ↪ by applying a phase on each superposition state
3. Use methods of interference to cancel or add phases to optimize for the correct answer
   ↪ and shrink the wrong answers (like noise canceling in headphones)

Challenges

- As more number of qubits are added to a system,
  ↪ ... the higher the error rates.
- stay free of outside interference (**decoherence**)
  ↪ ... for **as long as** it takes to finish a computation.

Quantum Software & Hardware (2018)

QAT (Quantum Application Toolset), Atos

Project Q

QisKit (IBM)

Q# (Microsoft)

Cirq (google)

XACC (oak ridge)

QS-Intel (Intel)

Forest (Rigetti)

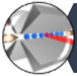

Trapped Ions
14 qbits
IonQ, innsbruck

Superconductor qubits
up to 72 qbits
IBM, Google, Rigetti

Spin qubits
Intel
49 qbits

Topological qubits
Microsoft

Courtesy of Dr. Jean-Noël Quintin, Quantum Computing Lab, Atos/Bull

Summary

Conclusion

- **Luxembourg government priority on HPC**
  - ↪ sustained by University of Luxembourg HPC developments
  - ↪ consolidate and extend Europe efforts on HPC/Big Data

UL HPC (as of 2018)

1029.342 TFlops / 9852.4TB (shared)

Conclusion

- **Luxembourg government priority on HPC**
  - ↪ sustained by University of Luxembourg HPC developments
  - ↪ consolidate and extend Europe efforts on HPC/Big Data

UL HPC (as of 2018)

1029.342 TFlops / 9852.4TB (shared)

Several On-going Strategic HPC efforts in Europe...

- **. . . in which UL *(HPC)* is involved . . .**
  - ↪ ETP4HPC, EU COST Action NESUS etc.
  - ↪ **PRACE** - Official representative for Luxembourg from UL
  - ↪ EuroHPC / IPCEI on HPC and Big Data (BD) Applications
  - ↪ **National HPC-BD Competence Center**
  - ↪ NVidia Cooperation agreement on AI and HPC

Quantum Computing Take away

- Relies on **Quantum mechanics**: laws of subatomic particles
  - ↪ super powers: superposition, quantum tunnelling, entanglement.
- Quantum computers use **qubits** to encode quantum information
  - ↪ calculate complex mathematical problems
  - ↪ can solve optimization problems, can revolutionize many fields
    - ✓ security/cryprography, drug discovery, AI etc.
  - ↪ unlikely that quantum computer will replace classical computers
- Qubits are **unstable**, very susceptible to environmental changes,
  - ↪ only work with free of outside interference due to fragile system
- Many types of qubits being used today

Quantum Computing Take away

- Relies on **Quantum mechanics**: laws of subatomic particles
  - ↪ super powers: superposition, quantum tunnelling, entanglement.
- Quantum computers use **qubits** to encode quantum information
  - ↪ calculate complex mathematical problems
  - ↪ can solve optimization problems, can revolutionize many fields
    - ✓ security/cryprography, drug discovery, AI etc.
  - ↪ unlikely that quantum computer will replace classical computers
- Qubits are **unstable**, very susceptible to environmental changes,
  - ↪ only work with free of outside interference due to fragile system
- Many types of qubits being used today

Lots of research still needed to create a quantum computer
... that can defeat our classical supercomputers

Thank you for your attention...

Questions?

http://hpc.uni.lu

High Performance Computing @ uni.lu

Prof. Pascal Bouvry
Dr. Sebastien Varrette
Valentin Plugaru
Sarah Peter
Hyacinthe Cartiaux
Clement Parisot

University of Luxembourg, Belval Campus:
Maison du Nombre, 4th floor
2, avenue de l'Université
L-4365 Esch-sur-Alzette
mail: hpc@uni.lu